Transparent Clouds: An Enhancement to Abstraction

Reza FARRAHI MOGHADDAM[1,2]
[1]Synchromedia Lab and CIRROD
ETS (University of Quebec)
Montreal, QC, Canada H3C 1K3
Email: imriss@ieee.org
LinkedIn: https://www.linkedin.com/in/rezafm

Yves LEMIEUX[2]
[2]Ericsson Research - Cloud Technology
Ericsson Canada Inc
Montreal, QC, Canada H4P 2N2

Mohamed CHERIET[3]
[3]Synchromedia Lab and CIRROD
Production Automation Department
ETS (University of Quebec)
Montreal, QC, Canada H3C 1K3

arXiv:1609.00664v1 [cs.NI] 2 Sep 2016

***Abstract*—With the introduction of various hardware/software technologies such as Cloud Technologies or Virtualization technologies, there has been a great potential to reuse ICT artifacts thanks to Abstraction and also Exchangeability features achieved via these technologies. These technologies also provide various advantages with respect to sustainability including resource consumption reduction (in the use phase only or in the whole life cycle). However, there is an additional but untapped potential associated with the anonymization of resources introduced by both abstraction and exchangeability features. By realizing on this potential, we can improve cloud solutions and reduce their by-product opacity, which usually prevents leveraging on the specialized but tweakable (i.e., nonessential modifications without changing the main function) features of components that are captured in the component models. This is especially a challenge in the case heterogeneous/disaggregated infrastructure where developing models to cover everything is practically impossible. In this work, by leveraging on the concept of pathways, we develop a few mechanisms that enable transparency and therefore tweakability of features even in the presence of abstraction and heterogeneity. In particular, the layered-stack approach to system decomposition is considered because of its role in both software defined networking (SDN) and Network Function Virtualization (NFV) system decompositions. For a concrete example, the case of dynamic frequency scaling of processors is considered and it is shown that the associated consumption could be considerably reduced without requiring additional changes to the middle components.**

I. Introduction

In the context of sustainability, the life cycle of material, products, and also *things* in general has been the one of important focus areas [1], [2], [3], [4], [5], [6], [7], [8], [9], especially in terms of reducing resource consumption. In the domain of the Information and Communication Technology (ICT)[1] there are various technologies that pave the way for better use of ICT artifacts available in each ICT solutions. Cloud Technology and Virtualization Technology are two of these enabling technologies, and they have been well recognized and used toward sustainability goals in ICT operations [10].

In Appendix 1, we will discuss both Cloud and Virtualization technologies in a greater details. Nonetheless, we can mention that there are two features that are the core of these technologies: i) Abstraction and ii) Exchangeability. Similar to many other use cases of these features [11], they have led to the introduction of opacity in the operations. Therefore, there has been a well-acknowledged compromise in association with the benefits of these features: Reduced adjustability as a result of opacity.

The paper is organized as follows. In Section II, the problem statement related to heterogeneous and disaggregated infrastructure is discussed. Then, in Section III, the proposed Northwise/Southwise Visibility/Transparency Pathway (NSVTP) approach is discussed, and it is followed with some special cases of targeted visibility. The use case of NSVTP in dynamic frequency scaling of processor cores toward reduction of energy consumption, increase of the equipment life time, and ultimately sustainability is provided in Section IV. Finally, the conclusions and the future prospects are presented in Section V. Also, in Appendix 1, a generic definition of cloud and virtualization technologies is provided. Then, the layered and networked stacks approach to service decomposition is presented in Appendix 2.

II. Transparent Cloud

Considering the definition of the cloud computing [12], which requires abstraction of components as the sixth key characteristics (Resource pooling), we define the *transparent cloud computing* as an enhancement to standard cloud computing in which the components expose their feature/characteristics as much as possible to each other. In many cases, the service providing components in a pair of connected components sits on underlying layers of the cloud stack. Therefore, these components, which we call them southern components, would be required to provide as much as possible information and control of their features to the service receiving component in the pair (the northern components in this analogy).

In the context of 'homogeneous' infrastructure, the abstracted model of a component would cover almost all the information and controls except the name (ID) of the component. The name (ID) of the components are by definition replaced with the virtualized names (IDs). This means that the cloud computing solution built on top of *homogeneous infrastructure* are by default transparent.

[1] or more practically (E)ICT that also includes the embedded forms

A. Heterogeneous Infrastructure and Transparent Cloud

In contrast to homogeneous infrastructure, the concept of 'transparent cloud' requires special attention in the context of the emerging heterogeneous/disaggregated infrastructure (especially hardware part). With the emergence of many new manufacturing/designing players in the hardware domain, the solutions (such as cloud) have no choice except to operate with a highly 'heterogeneous' set of hardware components. Even in the case of only a single manufacture, the set of possible hardware components is growing mainly in response to the diverse compute needs of the new applications. On top of that heterogeneous set of hardware components, the number of the possible ways to combine these hardware components is much bigger. It has been observed that a heterogeneous/disaggregated infrastructure is of great advantages for both the provider and the customer in terms of reducing the total cost and also overprovisioning [13].

It has been suggested to use a combination of managers and orchestrators in order to handle and operate a multi-tenant compute center (such as that of the case of virtualized network functions [14]). However, development of a true cloud solution that has the capability to take the maximum advantage from a heterogeneous infrastructure is a challenge mainly because of the abstraction aspect of the cloud solutions. A particular case of this challenge is discussed in the next section.

B. The Challenge of Resource/Component Rotation in Heterogeneous Infrastructure

A particular case in this context is the case of component 'rotation', which is an advantage of the cloud solution. In a component rotation action, an underlying component (which probably has failed) is replaced with another component with the same features and in the same layer without requiring any changes to the overlying components. This is thanks to the inherent abstraction feature of the cloud solutions.

However, in the case of heterogeneous infrastructure, the cloud operator may encounter cases in which they cannot find an available, replacing component with the same features, and therefore they are forced to use a component with 'higher' grade features in order to complete the rotation action. As mentioned before, avoidance of overprovision is a key advantage of heterogeneous infrastructure, and therefore the cases when a high-end component is used in place of a low-end component would be non-zero. Although the solution will work thanks to abstraction, the extra features of the high-end component will not be used. If we can find a way to expose the extra features of the replacing component to the northern component, not only higher performance could be achieved, the user experience (in this case, the tenant experience) would be increased by becoming aware of the free upgrade they received.

It is worth noting that there are no changes to the main service interactions among the components. In other words, the additional mechanism does not require any modifications on the connections in the cloud stack. Instead, it provides an extra means of communications between two components in order to adjust their states according to the preference of the other component even if that change in the state has not been captured by the abstracted model of the component. An example of such states is the clock frequency of CPUs. Changing the state of a CPU with respect to its clock frequency (using the particular features provided by the manufactures) would not require any change in the main stream of instructions being executed. However, such a change could be of the interest of the application for many reasons including cost saving, footprint reduction, and even demand response.

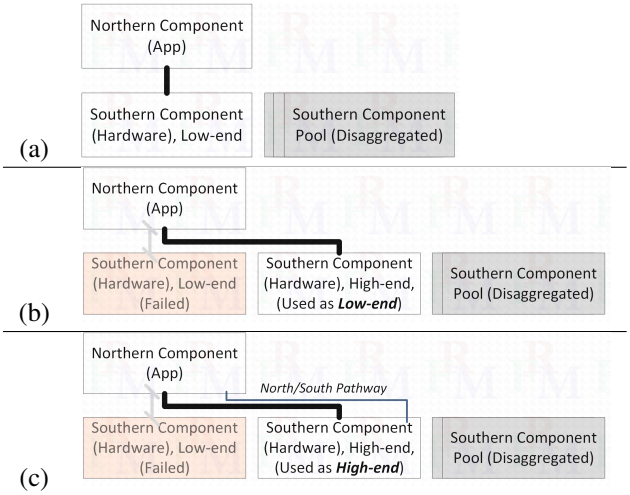

Fig. 1: A typical example of the application of NSVTP in a heterogeneous/disaggregated infrastructure environment. a) The requested setup. b) The configuration after failure of the southern component. A high-end (not required in the requested setup) is pulled from the pool of disaggregated hardware and used as a low-end component (low utilization). c) The same as (b) but with the NSVTP pathways enabling the northern component to adjust/modify the high-end features of the new southern component (high utilization possibly with extra billing). No other modification is needed.

Figure 1 shows a typical use case in which a southern component fails and is replaced by a high-end component. Two alternatives are considered. In the first one, there is no additional mechanism to expose the features of the new southern component to the northern counterpart. In the second alternative, which is the case of this paper, a direct messaging mechanism allows the exposure, and then consequently the northern component can adjust the state of the northern one to its preferred state in order to reduce the cost or improve the performance. In the case of homogeneous hardware, there will be no significant advantage for the second case, and both cases could be considered the same in terms of transparency.

C. NSVTP Approach to enable Transparent Cloud

Instead of redefining the complete stack of a cloud solution, here, we proposed an alternative that acts as an 'connection' (but not within the stack) in order to enable components (hardware, software, and services) to expose their specific features/states to each other (especially those in the higher layers) without requiring any modifications to the 'middle'

components. This proposed solution is based on an additional obstructed messaging between two components. Therefore, it does not affect the main message/API calls that passes among the components and services in their normal operation. Instead, it allows a component to request another component to change its state in order to better fit with the requirements of the former one. To reduce its impact to as low as possible, the messages are exchanged in the form of component names. In other words, the names are generalized in order to make them capable to carry the exposure and control messages without requiring addition of any extra services. The details of the proposed approach, which is called Northwise/Southwise Visibility/Transparency Pathway (NSVTP), is provided in the next section. It is worth mentioning that the proposed NSVTP approach and its mechanism is not the only way to realize the concept of transparent cloud in the context of heterogeneous infrastructure.

III. Northwise/Southwise Pathways Approach to Transparency and Visibility

The purpose of the Northwise/Southwise Visibility/Transparency Pathways (NSVTP) feature is to enable direct, immediate interaction and communication between two components in a stack while having a 'minimal' impact or footprint on the functions and APIs of any middle-layer component involved in the operation. To achieve such functionality, we break down the NSVTP in the two pathways: One to carry communications to the 'north' and one to handle the communications to the 'south'. It is worth mentioning that here we use the notions of north and south with respect to vertical dimension of the stack, and they should not be mistaken with the similar notions used in the context of chipset bridges.

Along each one of the two pathways, i.e., the northwise and southwise pathways, the NSVTP interactions are implemented in the form of some appendices, which we call 'capsules', that are added to the component IDs. Considering the fact that almost all currently available stack designs use a name-based approach to enabling interactions among their components, we assume that every component is assigned with a unique component ID (i.e., resource ID). In the following sections, we discuss how the NSVTP approach uses and also extends the name IDs in order to enable minimal-footprint insertion of direct interaction between not-immediately-connected components in a stack. Again, it is worth mentioning that these additional, inserted interactions/communications should not be mistaken with the actual interactions/communications between the components. The actual solution is carried forward using the actual interactions, while the additional NSVTP interactions enable modification of nonessential features of the (southern) components.

A. Northwise Blueprint/Status Capsules

The core of the Northwise capsules is an extension to the 'component' ('resource') identifier.[2] The advantage of this form of direct communication and interaction between two components is that there is no additional requirements or functions that the 'middle' components would be required to provide. In other words, the two components directly 'call' each other, and the middle components just pass the names along the way up (north) or down (south). In the proposed extension to the resource identifier, shown in Figure 2, a regular resource identifier (called a resource ID from here on) is augmented with an *appendix* that carries the capsule data. In many of the NSVTP messages or communications, the appendix could be omitted if the required data has been already communicated between the two components (and has been constant afterward). Figure 2 provides a basic structure of the extended identifier. It starts with the resource ID, i.e., the 'regular' resource identifier assigned to the component, and then it carries the capsule header followed by the capsule's actual data. The capsule header is used to differentiate among various types of capsules and also other possible appendices. The capsule data itself is divided into two subparts or segments: i) the blueprint and ii) the status. The blueprint segment is the core of the capsule, as described in Figure 3, in which a typical blueprint segment is shown. A blueprint consists of a header followed by a plurality of blueprint's *schemes*. Each scheme itself, as shown in Figure 4, consists of its associated parameters and *formula*. We assume that each scheme could have more than one formula in order to account for various *regimes* of operation of components without over-fitting or

[2] The extension is considered as an option.

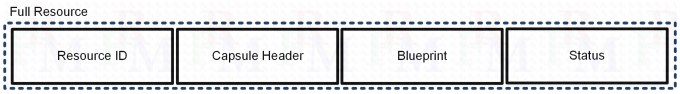

Fig. 2: The structure of a typical extended name ID including the northwise capsule appendix. The subparts of the capsule, i.e., the header, the blueprint, and the status segments, are also shown.

Fig. 3: The structure of a typical blueprint segment of a northwise capsule. A header followed by a plurality of schemes constitute the blueprint segment.

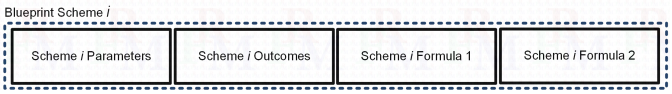

Fig. 4: The structure of a typical blueprint scheme. It consists of the parameters, the outcomes, and a plurality of formula.

under-fitting a formula in order to cover all the regimes.[3] As an example, the scheme of the processor cores with respect to their dynamic frequency is describe in Figure 8.

The status segment of the capsule data is used to transfer or inquire about the status of a southern component or the status of its constituents. The status data is populated using the parameters followed by the values in a simple to interpret way such as that of JSON data structures. In order to be more formalized, a scheme in the blueprint is reserved to formulate and provide the parameters and actions related to status segment.

By default, it is assumed that the capsule segments or the whole capsule itself are encoded (for example compression encoding for a reduced size). This allows for lowering the cost and communication load associated with the capsules transfer and also messaging. The encoding can be also replaced or enhanced with encryption in those cases where the capsule data is assumed to could potentially 'interfere' with the operation and performance of the 'middle-layer' components. However, it is also optionally possible to skip/omit encoding or encryption of the capsules at all. It is worth mentioning that the whole capsule can be enveloped in a Type-Length-Value (TLV) element as an optional information in order to be transparent to the involved protocols across the stack [15]. Also, there is no fixed space/size reserved for the segments in the capsule structure in order to save space when a segment is constant compared to previous communicated instances of an extended resource ID. When a segment is constant with respect to its previous description, it is replaced with NULL characters from the "alphabet" used to describe the resources or capsules. Also, by default, the blueprint and the status segments are encoded/encrypted separately from each other and then are concatenated. However, they could be first concatenated and then are encoded as a single entity, if preferred (or because of possible difference in the encoding performance in the actual setup).

B. Southwise Tweak Capsules

In the southwise communications, the capsule structure is similar to that of the northwise communications mentioned in the previous section. The main difference is that the appendix is composed of a capsule header followed with a plurality of *tweaks*. A 'tweak' is defined as an *populated* version of one of the formula from the associated blueprint schemes previously communicated in a northwise capsule to the northern component. By 'populated' we mean that the parameters of the scheme are set by the northern component.

[3] Every component could have different behaviors when it is subject to different input. For example, a component could have different *regimes* of operation at its nominal, over-nominal, or under-nominal input (activity) levels. If a model is forced to cover all these regimes, it is usually prune to under-fitting, i.e., lower but constant accuracy. In contrast, if a very small interval is used in the modeling, the model would be an over-fit, i.e., claiming a high accuracy but with the risk of unbounded errors outside the modeled interval.

C. Trusted eXchange (TX)

In cases where an encoded/encrypted NSVTP interactions is preferred, a trusted component is considered to handle the key exchanges and also verify the layer index of the components. The southern component s encapsulates an instance of its blueprint B_s in a capsule C_s using encoding/encrypting approaches. Then, it again encapsulates the capsule C_s and the associated key in another capsule C'_s probably using another approach (i.e., k'). Then, it sends *up* the second key k' to middle layers and finally to the *targeted* northern component on the l^{th} layer, and at the same time sends the tuple (C'_s, l) to the Trusted eXchange (TX), where l represents the layer *index* of the northern component. The northern component is assumed to contact the TX with presenting a tuple consisting of i) the key received, ii) its layer index, and iii) a key for southwise capsules (k''). We assume that the TX is capable to verify the *layer index* of the incoming requests. Then, the TX uses the key provided by the requester, i.e., the northern component, and retrieves C_s and its associated key, and then sends them to the requester.

In many cases, the TX could be practically omitted. Then, direct exchange of blueprints and tweaks between the southern component and northern component is used. The two parties could still use encoded/encrypted copies in the exchanges in order to reduce the size of messages being exchanged. However, verification of the layer index would be a challenge if required.

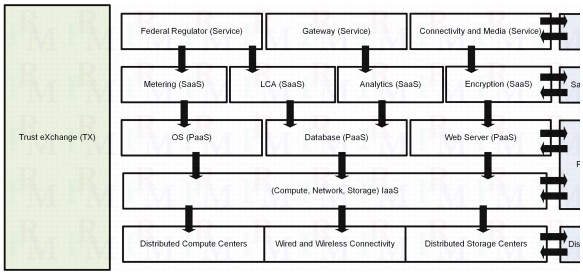

Fig. 5: Layered Service Stack along with its Trust eXchange.

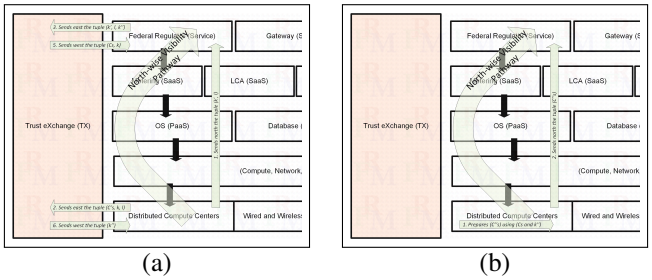

(a) (b)

Fig. 6: The time steps of a northwise path establishment using the Trust eXchange in a Layered Service Stack. a) The initiation steps. b) Northwise communication after the initiation steps.

The actual step to perform and initiate a NSVTP between a southern and a northern components is shown in Figure

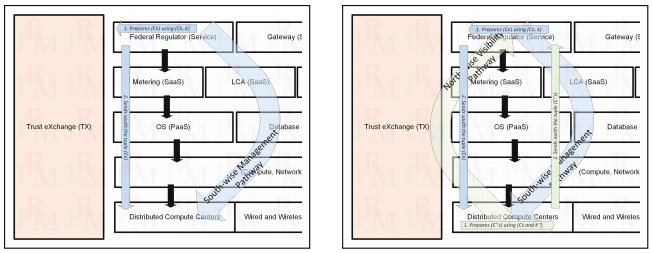

(a) (b)

Fig. 7: a) The time steps of a southwise path establishment using the Trust eXchange in a Layered Service Stack. b) The overall north/south interaction after initiation step (no more need for the TX mediation).

6(a). After the initiation steps, the southern component could directly communicate and call the northern component, as shown in Figure 6(b). At the same time, the northern component directly calls the southern component after the initiation process (Figure 7(a)). The complete bi-directional NSVTP interaction between the two components is shown in Figure 6(b).

IV. Use Case of Dynamic Frequency Scheme for Processor Cores in a Layered Stack

Power consumption and performance modeling and management in data centers is of great interest [16], [17], [18], [19], [20], [21]. The Dynamic Frequency Management of processor/cores' clock frequency has been considered and used as a alternative option to consolidation in saving energy and also in increasing the life time of equipment [8], [9], [22]. This approach is usually referred to as Dynamic Voltage and Frequency Scaling (DVFS) or other similar terms that actually refer to those processes that are used to dynamically change the clock frequency. However, from here on, we use the term DVFS in a generic form to cover all realization processes for the purpose of simplicity. It has been shown that DVFS-enabled, workload-aware managers in the data centers and clouds could highly reduce the OPEX of the operation thanks to saving in the energy consumption by the processors [9]. This saving is usually rooted in a 'nonlinear' relation between the electricity power requirement of a core and its clock frequency. In its simplistic form, a cubic relation could be inferred. However, in the real cases, the power-law index is usually less than 3:

$$P = (P_0 + P_3(f/f_{\max})^{n_{\text{DVFS}}}(l/l_{\max})$$

where P is the power required by the core, f is the current clock frequency, $f_{\max}$ is the maximum clock frequency possible, l is the load on the core in the Million Instructions per Second (MIPS) unit, and n_{DVFS} is the associated power-law index.

Also, it can be argued that many applications, such as Map-Reduce, follow a form of deterministic cycles of 'isolated intense compute' followed by 'inter-node intense transfer' [23]. Although proposing new patterns for application could be a necessity, optimizing the resource and application based on the current patterns is another requirements for the cloud manager/controller and in general for the Data Center Infrastructure Management (DCIM). In what follows, we provide an example of how NSVTP approach could be used in order to leverage on such patterns in applications.

The adjustment of DVFS by the manager/controller of the cloud system is great. However, it could be 'skipped' by the application developers who could see it as a risk factor for their performance. At the same time, the un-guaranteed latency introduced because of the communications required to the manager/controller and then from there to the southern components makes it practically impossible to perform the DVFS adjustments in very 'fine' granularities along the time. As mentioned before, there is a great potential for optimization and energy saving in fine granularities because of periodic patterns of compute/communicate cycles in many applications. This brings us to the proposed NSVTP approach which enables real-time, direct communication between a northern component (probably an application) and a southern component (such as a CPU core) in order to fully utilize the benefits of dynamic adjustments in the clock frequency.

Figure 8 shows an example of blueprint scheme that could be probably provided by a southern CPU core to a northern component in order to set up a DVFS 'tweak'. There are two parameters, the frequency step and the latency. Each parameter has a set of feasible values to select from. For example, the latency parameter has only one value (t_1) in its feasible set which means each DVFS interactions from the northern component would take at least a time interval of t_1 in order to change the frequency. That means if the compute/commute (or intense/relax) pattern of the application has a time period around or smaller than t_1, the dynamic frequency scaling would not have any benefit to the operation. This shows the capability of the NSVTP in enabling DVFS at the time scales as small as the latency limits of the DVFS changes. This potential can be leveraged toward energy saving, increase in the performance, and in general sustainability. Considering the norm that the DCIMs usually modify the state in a long period of time (such as one hour), it is clear that the NSVTP operating at short time intervals (such as one second or less) could bring in semi-autonomous, decentralized management mechanisms without requiring any modification neither in the middle-layer components nor in DCIM components. In the following subsection, we provide some estimations of the impact of NSVTP addition to a cloud computing operation.

The latency in the dynamic frequency scaling interactions could range up to 20 µs [24]. However, it should be pointed out that this time interval could be actually a halt/pause in the processor/core operation during which the voltage and frequency is adjusted. In this work, we assume that this time interval is only a transition time interval in which the processor/core's circuit is active but not stable (such as that of the case of Marvell XScale processors [25]).

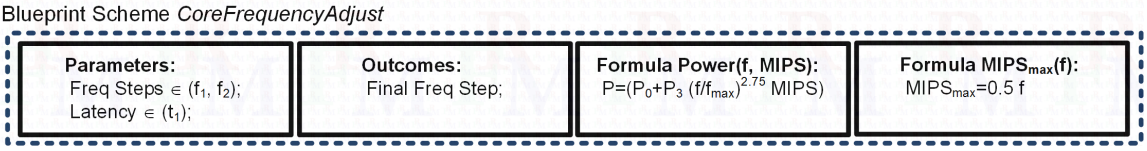

Fig. 8: A Blueprint Scheme related to DVFS use case.

A. Use Case of Dynamic Frequency Scheme: Experiments

Although each application could have a different configuration requirements or best practices, we assume that application of interest in our study of long-term with periodical compute-commute intervals. Also, we assume that the applications are hosted in a cloud-oriented data center with disaggregated hardware at the cabinet (rack) scale [26], [13]. That means that the hardware component available in a cabinet could be automatically mixed-and-matched in timely manner in order to build desired compute nodes up to the maximum capacity of the cabinet. It is worth mentioning that, even without NSVTP, disaggregated hardware provides a considerable amount of energy saving and reduction in footprint thanks to its ability to reduce over-provisioning and to "recycle" the hardware components.

In the experiments, we use the same specifications for the DVFS feature as those of [8]: $P_0 = 142.2$W and $P_3 = 107.8$W, $f_{\text{min}} = 1$ GHz, and $f_{\text{max}} = 3$ GHz. Also, the performance of the hardware (thanks to disaggregated granularity) is assumed to be 80%, i.e., on average for a workload of intensity 1 MIPS, we would allocate a 1.2 MIPS core equivalent because of discrete nature fo the number of cores. This overhead is because of discrete nature of the number of cores, and it is not specifically dependent on 1 MIPS. For example, for cores of a power of 3 MIPS, all workloads with a MIPS less than 3 MIPS would be assigned at least one core. Resource sharing practices would allow other workloads to use the extra compute power of a core. However, because of requirements such as isolation, limiting management overhead per resource, and handling unpredicted workload variations, the maximum performance would be limited below 80%.

In the first case, we assume only one application. Considering the disaggregated nature of the data center, this means that we are ignoring the rest of application in this case, and therefore there is no conflict with multi-tenancy operation of the data center. The cycles of compute-commute of an application i can be characterized by the ratio compute sub-interval to the data exchange sub-interval (ρ_i) and also the compute sub-interval itself ($t_{\text{comp},i}$). These two application characteristics could be modeled using statistical distributions. However, here we assume that they are represented by their mean values. The performance of the application without using NSVTP in terms of energy consumption could be calculated as follows: Over each cycle of compute-commute pattern, the associated processor/core would consume $(P_0 + P_3)\,(t_{\text{comp}})\,(1 + 1/\rho)$. When the NSVTP is used, assuming $\delta < 2 * t_{\text{comp}}/\rho$, the consumption could be calculated as follows: $(P_0 + P_3)\,(t_{\text{comp}} + 2\delta) + \left(P_0 + P_3\,(f_{\text{min}}/f_{\text{max}})^3\right)(t_{\text{comp}}/\rho - 2\delta)$. Therefore, the overall improvement because of NSVTP would be:

$$\eta = \frac{(1 + 2\delta/t_{\text{comp}})}{(1 + 1/\rho)} + \frac{\left(P_0 + P_3\,(f_{\text{min}}/f_{\text{max}})^3\right)}{(P_0 + P_3)} \frac{(1/\rho - 2\delta/t_{\text{comp}})}{(1 + 1/\rho)} \quad (1)$$

Figure 9 shows the impact of these parameters on the performance of the NSVTP approach in reducing the energy consumption of the application in the compute nodes. Because we do not optimize the behavior of the application in this work, the energy consumption and other impacts of the exchange and transfer of data sub-intervals would be the same (therefore, we ignore the modulation phenomena). As can be seen from the figure, the NSVTP mechanism provides an additional minimum of 60% relative improvement compared to the baseline operation, i.e., not using the NSVTP mechanism. The minimum improvement is related to high ρ value regions of operation that are related to those applications that have long period of compute-intense time intervals compared to their data-exchange time intervals. It is worth mentioning that the key factor in 'lower' performance (but still positive) of the NSVTP mechanism for high-ρ regions is the presence of disaggregated hardware in our use case. In other words, the disaggregated hardware feature (along with presence of a heterogeneous hardware resource pool) allows allocation of almost-optimal resources to applications with and without inclusion of the NSVTP mechanism. Therefore, for settings that are not using the disaggregated hardware feature in their cloud solution, the positive impact of the NSVTP mechanism would be much higher. We consider such cases in a future work.

B. Impact of N/S Pathways in Reduction of Footprint Associated with Widespread Adaptation of Cloud Computing and ICT

Considering the forecasts that the data center market will grow rapidly in the coming years (such as 22% Compound Annual Growth Rate (CAGR) forecast for the Nordic region market [27]), there will be a considerable potential for the NSVTP-based approaches and mechanisms to reduce the energy consumption of the data center and cloud operations which in turn improve the performance and also reduce the associated environmental footprint. With the minimum of 60% improvement obtained for the DVFS use case in the previous section, combined with the disaggregated hardware

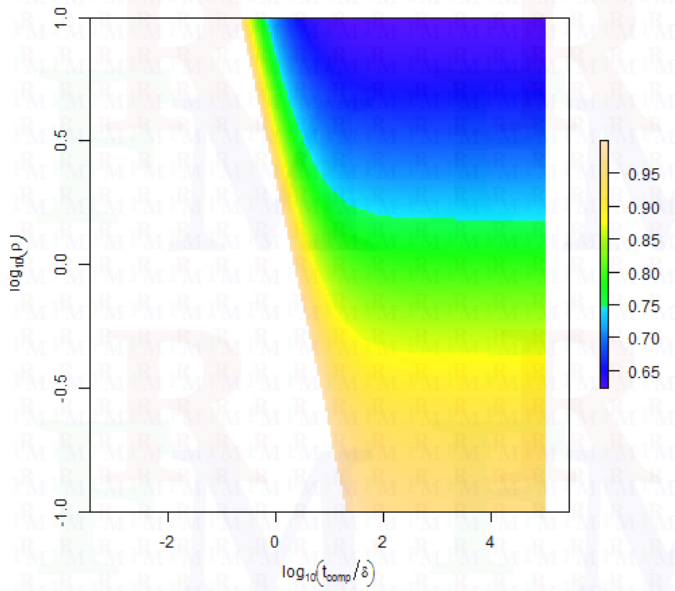

Fig. 9: The performance of NSVTP in the case of a single application. The performance is plotted versus two dimensions: The relative t_{comp}/δ and ρ. A performance value of for example 0.7 in the figure would indicate a 70% decrease in energy consumption compared to the baseline. The minimum value for the performance is around 60% in this case.

performance, the NSVTP approaches and other similar decentralized management and operation mechanisms would bring a considerable level of energy and cost saving, which in turn translate in lesser environmental footprint and higher level of sustainability.

V. Conclusion

Starting from the standard approaches to service stack decomposition, specifically layered stack decomposition and networked stack decomposition, a new noninvasive approach and mechanism to automation of cloud computing operation has been proposed. The proposed approach, called Northwise/Southwise Visibility/Transparency Pathways (NSVTPs), enables the northern components such as applications to directly communicate with southern components such as hardware without adding any additional requirement on the middle-layer components. The NSVTP provides this means of communication toward bypassing the opacity induced by the middle-layers and abstraction which have been well-accepted by-products of cloud computing. The core of the NSVTP approach is based on a concept called capsules that is then integrated with the resource identifiers in the form of appendices. Capsules enable exchange of blueprints, schemes, and formula between southern and northern components. These exchanges could be then leveraged toward a more optimal unified operation of all components. In a special use case related to dynamic frequency scaling and DVFS, it has been shown that the NSVTP could improve the 'relative' energy performance with a further minimum of 60% even when it is added to the highly-adaptive and optimized disaggregated hardware solutions.

Acknowledgment

The authors thank the NSERC of Canada for their financial support under Grant CRDPJ 424371-11 and under the Canada Research Chair in Sustainable Smart Eco-Cloud (NSERC-950-229052), and also the MITACS of Canada.

References

[1] S. Pfister and P. Bayer, "Monthly water stress: spatially and temporally explicit consumptive water footprint of global crop production," Journal of Cleaner Production, vol. 73, no. 0, pp. 52–62, Jun. 2014.

[2] J. Malmodin, D. Lundén, Å. Moberg, G. Andersson, and M. Nilsson, "Life cycle assessment of ICT," Journal of Industrial Ecology, vol. Online First, May 16 2014.

[3] F. Stiel and F. Teuteberg, "Measuring the environmental impact of IT/IS solutions - a life cycle impact modelling approach," Environmental Modelling & Software, vol. Online First, no. 0, pp. –, 2014.

[4] Y. Arushanyan, E. Ekener-Petersen, and G. Finnveden, "Lessons learned — review of LCAs for ICT products and services," Computers in Industry, vol. 65, no. 2, pp. 211–234, Feb. 2014.

[5] T. Schaubroeck, R. A. F. Alvarenga, K. Verheyen, B. Muys, and J. Dewulf, "Quantifying the environmental impact of an integrated human/industrial-natural system using life cycle assessment; a case study on a forest and wood processing chain," Environ. Sci. Technol., vol. 47, no. 23, pp. 13 578–13 586, 2013.

[6] D. Schien, P. Shabajee, M. Yearworth, and C. Preist, "Modeling and assessing variability in energy consumption during the use stage of online multimedia services," Journal of Industrial Ecology, vol. 17, no. 6, pp. 800–813, 2013.

[7] R. Farrahi Moghaddam, Y. Lemieux, and M. Cheriet, "The service-bond paradigm — potentials for a sustainable, ICT-enabled future," EAI Endorsed Transactions on Smart Cities, vol. 16, no. 1, p. e1(17), Jul 18 2016, [arXiv preprint http://arxiv.org/abs/1507.06295 arXiv:1507.06295, Jul 2015].

[8] R. Farrahi Moghaddam, F. Farrahi Moghaddam, T. Dandres, Y. Lemieux, R. Samson, and M. Cheriet, "Challenges and complexities in application of LCA approaches in the case of ICT for a sustainable future," in ICT4S'14, Stockholm, Stockholm, Sweden, Aug 24-27 2014, pp. 155–164, [ArXiv preprint: http://arxiv.org/abs/1403.2798 arXiv:1403.2798, March 2014].

[9] F. Farrahi Moghaddam, "Carbon-profit-aware job scheduling and load balancing in geographically distributed cloud for HPC and web applications," Ph.D. dissertation, École de technologie supérieure (ETS), University of Quebec (UduQ), Montreal, Quebec, Canada, Jan 16 2014.

[10] R. Farrahi Moghaddam, F. Farrahi Moghaddam, and M. Cheriet, "A multi-entity input output (MEIO) approach to sustainability - water-energy-GHG (WEG) footprint statements in use cases from Auto and Telco industries," 2014, [arXiv preprint http://arxiv.org/abs/1404.6227 arXiv:1404.6227, April 2014].

[11] A. Bourouis, K. Klai, Y. E. Touati, and N. B. Hadj-Alouane, "Opacity preserving abstraction for web services and their composition using SOGs," in ICWS'15, New York, NY, USA, Jun 27 - Jul 2 2015, pp. 313–320.

[12] ITU-T, "ITU-T Recommendation Y.3500: Information technology — Cloud computing — Overview and vocabulary," ITU, Geneva, Geneva, Switzerland, Series Y: Global Information Infrastructure, Internet Protocol Aspects and Next-generation Networks — Cloud Computing Edition 1.0, Aug 13 2014.

[13] Mainstay, "An economic study of the hyperscale data center," Mainstay, LLC, Castle Rock, CO, USA, White Paper, Jan 2016, [Available online: http://cloudpages.ericsson.com/transforming-the-economics-of-data-center].

[14] K. Bogineni, D. Davidson, A. Slauson, L. Molocznik, C. Martin, K. McBean, D. Brayshaw, T. Cagenius, T. Khan, K. Cramer, K. Smith, J. Gasparakis, P. Busschbach, N. Cocker, J. Hongisto, S. Scarlett, R. Iontel, G. Keegan, R. Yuhanna, and N. Anter, "SDN-NFV reference architecture," Verizon, New York City, NY, USA, Verizon Network Infrastructure Planning Version 1.0, Feb 24 2016.

[15] C. Dearlove and T. Clausen, "TLV naming in the mobile ad hoc network (MANET) generalized packet/message format," Internet Engineering Task Force (IETF), RFC 7631, Sep 2015, [Available online: https://tools.ietf.org/html/rfc7631 IETF:RFC4655].

[16] L. Piga, G. F. Gomes, R. Auler, B. Rosa, S. Rigo, and E. Borin, "Assessing computer performance with stocs," in ICPE'13, Prague, Prague, Czech Republic, Apr 21-24 2013, pp. 315–318.

[17] L. Bulej, T. Bureš, V. Horký, J. Kotrč, L. Marek, T. Trojánek, and P. Tůma, "Unit testing performance with stochastic performance logic," Automated Software Engineering, vol. First online on Jan 13, 2016, pp. 1–49, 2016.

[18] T. Bures, P. Hnetynka, J. Kofron, R. Al Ali, and D. Skoda, "Statistical approach to architecture modes in smart cyber physical systems," Charles University, Prague, Prague, Czech Republic, Department of Distributed and Dependable Systems D3S-TR-2015-04, Dec 22 2015.

[19] V. M. Weaver, "VMW group power measurement techniques," University of Maine, Orono, ME, USA, VMW Group UMAINE-VMW-TR-POWER-MEASURE-2015-08, Aug 26 2015.

[20] I. Lopez, S. Moore, and V. M. Weaver, "A prototype sampling interface for PAPI," in XSEDE'15, St. Louis, MO, USA, Jul 26-30 2015, pp. 1–4.

[21] X. Wang, B. Zhao, T. Mak, M. Yang, Y. Jiang, and M. Daneshtalab, "On fine-grained runtime power budgeting for Networks-on-Chip systems," IEEE Transactions on Computers, vol. Online First, 2016.

[22] A. Mazouz, B. Pradelle, and W. Jalby, "Statistical validation methodology of CPU power probes," in Euro-Par'14. Porto, Porto District, Portugal: Springer, Aug 25-26 2014, vol. Revised Selected Papers, Part I, pp. 487–498.

[23] E. Zahavi, I. Keslassy, and A. Kolodny, "Distributed adaptive routing convergence to non-blocking DCN routing assignments," IEEE Journal on Selected Areas in Communications, vol. 32, pp. 88–101, 2014.

[24] J. Suh and M. Dubois, "Dynamic MIPS rate stabilization in out-of-order processors," SIGARCH Comput. Archit. News, vol. 37, no. 3, pp. 46–56, 2009, presented first at ISCA'09.

[25] N. Poole, Y. Mo, J. Brusey, M. Langley, and R. Hazelden, "Power consumption of microelectronic equipment for wireless sensor networks," in Nanotech. NSTI, 2009, vol. 1, ch. 6: Wireless MEMS, pp. 534–537.

[26] Ericsson Inc, "Next-generation data center infrastructure," Ericsson Inc, Stockholm, Stockholm, Sweden, White Paper Uen 284 23-3264, Feb 2015, [Available online: http://www.ericsson.com/res/docs/whitepapers/next-generation-data-centers.pdf].

[27] WGR, "Data center construction market in the Nordic region 2016-2020," Wise Guy Reports, Nagpur, Maharashtra, India, Tech. Rep. WGR390256, Apr 7 2016.

[28] J. Malmodin and P. Bergmark, "Exploring the effect of ICT solutions on GHG emissions in 2030," in EnviroInfo'15 & ICT4S'15, Copenhagen, Hovedstaden, Denmark, Sep 7-9 2015, pp. 37–46.

[29] 3GPP, "IP multimedia subsystem (IMS). stage 2 (3GPP TS 23.228/ETSI TS 123 228)," ETSI, Sophia Antipolis, Alpes-Côte d'Azur, France, Technical Specification version 12.8.0, Release 12, Mar 2015.

APPENDIX 1. CLOUD AND VIRTUALIZATION

There are well-established and agreed-on standards available for both cloud computing and virtualization. For example, we can refer to the ITU-T recommendation with respect to the cloud computing's overview and vocabulary [12]. These standards, recommendation and guidelines are especially designed and used in the context of service request/provide interactions and transactions between cloud computing providers and users/customers. However, we think that the core of these technologies could be described at even a lower level. We acknowledge that such low-level description may not be feasible to be mapped to the provider/user interactions.

A. Cloud

Cloud Computing could be imagined as a plurality of computing operations on top of some 'nameless' resources. Although pluralized resources could be used to create a state of nameless resources, it is not the only possible way. In other words, we could reach the core of the Cloud Computing approach, i.e., relaxation of the dedicated-resource constraint, without enforcing plurality to the resources.

In terms of all efforts related to Cloud Computing standards, it could be mentioned that their coverage and convergence has been well carried out which is a key requirement for handling business operations. However, by nature, this has introduced a 'black-box' perspective that in turn masks the actual internal technologies and mechanisms used in each solution that granted the title of Cloud Computing. For the purpose of having an alternative 'white-box' perspective, it would be helpful to call one of the possible classes of internal approaches the Cloud Approach. This specific class would cover all those solutions which make resource 'names' (IDs) impermanent in order to allow 'resources' to be seamlessly changed (swapped, added, or dropped) by the user or the cloud manager. This class would include many possible approaches including those that stand on top of the Virtualization technologies. The virtual resources generated by virtualization approaches are by nature ephemeral and therefore their names carry the same ephemerality feature.

B. Virtualization

Virtualization could be seen as a way to make a resource plural. In other words, this technology provides a way to have isolated resource-sharing. This makes virtualization a good choice to separate and also to abstract the actual [physical] resources, which in turn makes virtualization a widely-used approach to enable cloud computing (with replacing the dedicated-resource constraint with resource pooling from a large set of short-living, ephemeral virtual resources). However, as mentioned in the previous section, it is possible to deliver a cloud computing solution without using virtualization.

To provide a metaphoric example, let us consider the case of vehicles. A 'Vehicle' could be defined a thing that provides the transport 'function' by relaxing the fixed-position constraints of some of its parts. In contrast, 'Wheel'-ization enables a disk object [resource] to have plural states (different angles) around its axis. Many vehicular types are delivered using the wheel technology. However, the question is that is it possible to have a vehicle without using wheels? And the answer in this metaphoric example is yes: For example, many construction machines, which are also vehicles, do not use wheels and instead they may use Continuous Tracks. However, the question would be then whether we can use these low-level definition of a vehicle in the market for provider/user interactions? The answer is simply no. Depending on jurisdiction, there are many standards and conditions that a vehicle should satisfy to be eligible to enter the market and the roads. For example, all vehicles should comply to Motor Vehicle Safety Regulations

(C.R.C., c. 1038) in Canada. The some logic would hold for the cloud technology. However, the regulatory guidelines and standards should not be literally taken as the 'boundary' of a technology by the developers: With introduction of any new instance or new feature to the cloud technology, there will be a period time of convergence needed to include that new instance/feature in the market regulations.

In the following section, we introduce a new feature in the cloud computing that could bring a higher level of performance and therefore greenness and sustainability to the operations if other influential factors are set properly. This feature, which we call the Northwise/Southwise Visibility/Transparency Pathways (NSVTPs), enables direct, real-time interaction among components of a cloud computing solution even if they are opaqued and blocked by the other middle-layer components. In order to put the NSVTP mechanisms in a proper context, we first define two major approaches to service decomposition in the cloud computing. Any effort to improve the overall sustainability and reduce resource consumption can be seen as an enhancement to cloud solutions. One category of such efforts can be modeled in the form modifying paths of operations into 'Re-Cycle' *paths*. These paths could be then augmented by other alternatives such as Up-Cycle modifications. Generally, we could categorize these modified paths in a bigger class of *X-Cycle* modifications. It is worth mentioning that such enhancements have had also an indirect positive consequence in the role of ICT as an enabler of sustainability in the non-ICT operations [28].

APPENDIX 2. LAYERED STACKS VS NETWORKED STACKS: NFV-NATIVE ARCHITECTURES

Although the Layered (in either Single-Column or Multi-Column forms) approach to the design of computing *stacks* has been generally favored, new approaches such as that of the Networked Stacks, which will be discussed in section Appendix 2-B, have been also considered mainly because of their specific advantages especially in the context of the 'function'-oriented architectures [7], [14]. Below, we first review the definitions of the layered and networked stacks, and then we will explore how it would be possible to add the proposed *pathways* to these types of stacks in the next section.

A. Layered Stacks [Approach to Service Decomposition]

A Layered Service-Stack Decomposition is a realization of a possibly-complex service in terms of an 'layer-indexed' set of a few simpler services/components. The simpler services are constrained to *only* interact with those other simple services that are their immediate neighbors in terms of the *layer* order. For example, a service with a layer *index* of the number 5 could only interact with services with the layer index numbers 4 and 6.[4][5] Any layered service decomposition could be visualized as a stack of simple services placed on top of each other according on their layer-order (for example, starting from the components with the *lowest* order index). In such a visualization, each simple component is considered as a service-layer component, and the whole picture is called a layered service-stack (see Figure 10).

B. Networked Stacks [Approach to Service Decomposition]

In contrast to the Layered Service-Stack Decomposition approach (and also the Chained Service-Stack Decomposition approach [7] which is not mentioned here), in the networked Service-Stack Decomposition there is no constraint in terms of the order and also in terms of the number of egress connectors exiting from a component (see Figure 11). In other words, any simple service could egress to more than one other simple services. A direct consequence of this relaxation of constraints would show itself as a challenge in the service orchestration, i.e., the 'timing' of service requests among simple services. At the same time, networked stacks have a great potentials in translating complex services in the form of functional relations among a plurality of simple components. The function-oriented approach has been of great interest especially in the context of Network Function Virtualization (NFV) toward less over-provisioned, more efficient and performance, and therefore more sustainable ICT operations. It is worth mentioning that the networked Stacks could partially support layer-order in some cases. We call such stacks quasi-linear networked stacks. An example is shown in Figure 11(a). In this figure, each cluster of Virtualized Network Functions (VNFs) stands on top of a thin layer of software defined infrastructure (SDI). The concept of pathways, which will be discussed in the next section for the layered stack decompositions, could directly applied within each 'island' of VNFs. The generalization of the pathways approach to inter-island interactions, which will discussed in the future, is of great interest when the networked stack decompositions move outside the quasi-linear form. In the unconstrained form, shown for example in Figure 11(b), there is no horizontal concept of location applicable within each island. However, still the vertical dimension along the functionalities is applicable and is used in the pathway mechanisms. To populate the VNFs in the figure, we used a few sample functions from the IP Multimedia Subsystem (IMS) solutions [29], [14]: Home Subscriber Service (HSS), Border Gateway Function (C-BGF), Resource and Admission Control Sub-System (RACS), Proxy-Call Session Control Function (P-CSCF), Serving-Call Session Control Function (S-CSCF), WebRTC[6] Web Server Function (WWSF), and Teleport (TP). Here, we added the TP function which allows seamless function-calling among VNF islands.

[4] It is worth mentioning that there could be exceptions in which a layer extends across more *than* one layer, and therefore could potentially interact with components that are at two different but adjunct layer index.

[5] Please note that the layer 'index' is not limited to be only *numeric*. For example, the set of {Infrastructure, Platform, Software, Service} could serve as the set of layer indices. However, we require the existence of an *order* among indices.

[6] WebRTC: Web Real-Time Communication

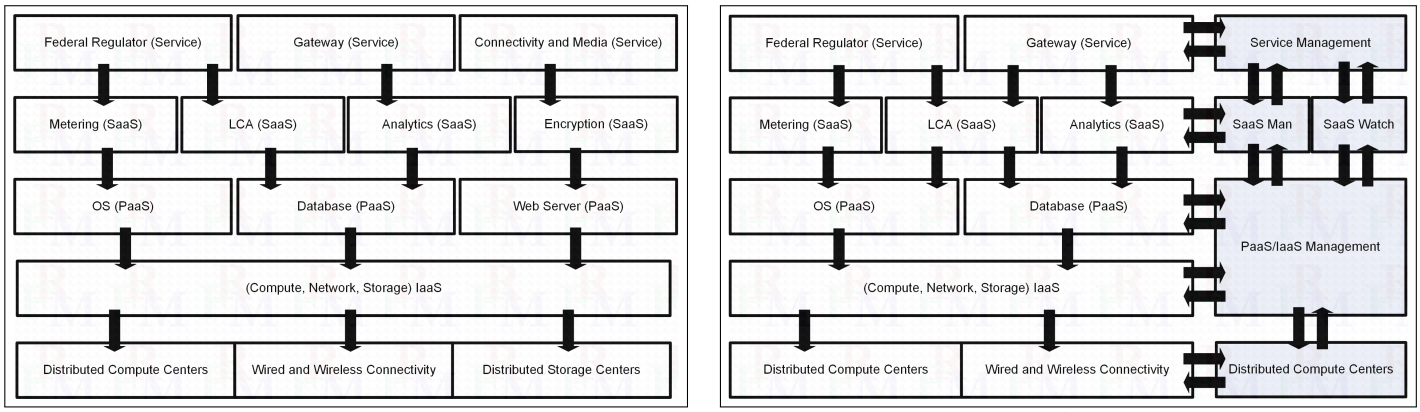

Fig. 10: Layered Service Stack. a) A Single-Column Layered Stack. b) A Multi-Column Layered Stack.

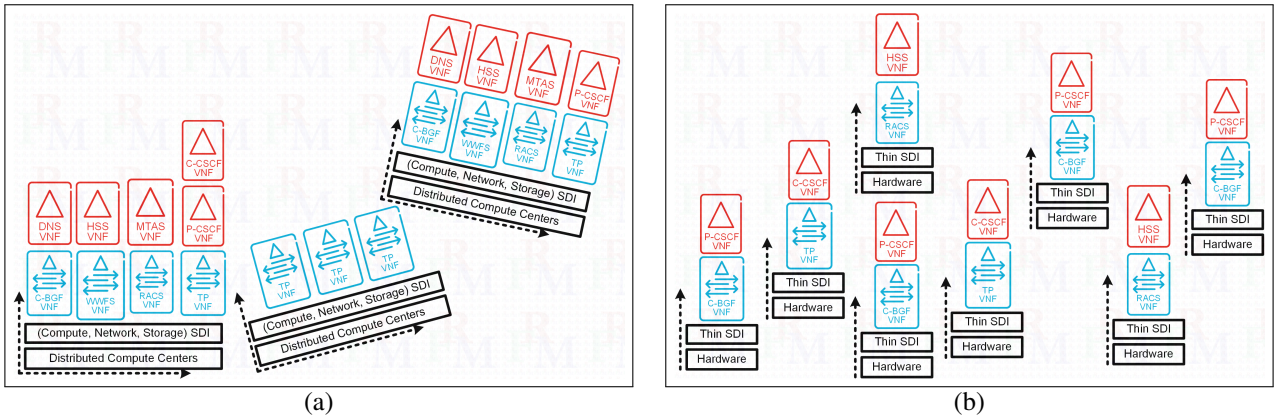

Fig. 11: Networked Service Stack. a) Quasi-Linear networked Stack. The horizontal axis indicates the 'location' and the vertical axis refers to the 'functions'. b) Unconstrained Networked Stack. There is no significant role for the location within the small islands. However, the functions axis is always relevant.